arXiv:1410.1501v1 [math.GT] 6 Oct 2014

When wild singularities of translation surfaces imply infinite genus

Anja Randecker

Karlsruhe Institute of Technology
Karlsruhe, Germany
e-mail: anja.randecker@kit.edu

We consider translation surfaces for which the metric completion is not necessarily a compact surface. Certain conditions on the additional points in the metric completion and on the dynamics of the translation surface imply that the surface has infinite genus.

Classical translation surfaces are objects at the intersection between many different fields such as dynamical systems, Teichmüller theory, algebraic geometry, topology, and geometric group theory. The history of translation surfaces starts in the time of the article [FK36]. Fox and Kershner obtained translation surfaces in the theory of billards when "unfolding" polygons with rational angles. The area became a field on its own with the seminal work of Veech in [Vee89].

The most visual description of classical translation surfaces is given by considering finitely many polygons in the plane. Every edge of the polygons is glued to a parallel edge of the same length so that we obtain a connected, orientable surface. The resulting object is locally flat at all points with the possible exception of the vertices of the polygons which will be called *singularities*.

A natural generalization is to drop the condition that the number of polygons has to be finite. When gluing infinitely many polygons, the local flatness still holds but the behaviour of the singularities is more diverse than in the classical case. This kind of translation surface is often called *infinite* in the literature, but we will not follow this convention here and simply call it *translation surface*.

Recently, the interest in this generalization of translation surfaces has grown: There are results on Veech groups in [Cha04], [HS10], and [PSV11], results on the dynamics in [Hoo14], [Tre14], and [LT14], and results on infinite coverings of finite translation surfaces (especially for the wind-tree model) in [DHL11], [HLT11], [AH12], [HW12], [HHW13], and [FU14]. However, while we have a classification of finite translation surfaces by studying strata of the moduli space, there is no systematic description for the generalized ones so far. A first step towards such a classification can be to understand and classify the singularities of the translation surfaces.

When considering translation surfaces with interesting singularities, it is natural to start with translation surfaces with exactly one singularity. So, a lot of the recently described examples are Loch Ness monsters, i.e. surfaces with infinite genus and one end (cf. [Ghy95] for the name of Loch Ness monsters and [Ric63] for the definition of ends and the classification of non-compact surfaces using ends). For instance, these examples include the baker's map surface in [Cha04], the Arnoux-Yoccoz surfaces in [Bow13], some examples from [Hoo10], and the stack of boxes in [Bow12].

Fundamental for this article is the work of Bowman and Valdez in [BV13] where they study *linear approaches* and *rotational components* of singularities. We will recall the definitions in section 1.

In this article, we want to present a link between the topology and the geometry of translation surfaces with so-called wild singularities. For this, a third aspect is used: the dynamics of translation surfaces. The main theorem we will show is:

Theorem 1
Let (X, ω) be a translation surface with the following properties:

(i) The singularities of (X, ω) are discrete.

(ii) There exists a wild singularity σ such that none rotational component of σ has length π.

(iii) There exist two directions θ_1, θ_2 for which the geodesic flow $F_{\theta_1}, F_{\theta_2}$ is recurrent for almost every point in X.

Then X has infinite genus.

For example, this theorem is applicable for all of the Loch Ness monster examples listed above (although for each of them it is easy to check by hand that the genus is infinite). Moreover, for the interesting class of parabolic surfaces (that is surfaces that have no Green's function, cf. the introduction in [Tre14]) the condition on the dynamics holds. So, to prove infinite genus of a parabolic translation surface using this theorem it is sufficient to check the conditions on the singularities.

The article is structured as follows: At first we give the definitions on translation surfaces and their singularities in section 1. In section 2, a criterion is developed how infinite genus can be shown using saddle connections. By establishing the existence of short saddle connections in section 3 we show that the conditions of the infinite genus criterion are fulfilled under the assumptions of Theorem 1. At this point, we will summarize the results and prove the theorem. In section 4, possible generalizations of the theorem are discussed by considering examples that also have interesting properties on their own.

Acknowledgements. The author wants to thank her advisor Gabi Weitze-Schmithüsen for hundreds of valuable discussions and her constant encouragement, Pat Hooper for many helpful discussions (especially at the beginning of this project, cf. Example 25),

and Ferrán Valdez for sharing his thoughts on singularities of translation surfaces, starting with explanations on [BV13]. Furthermore, the author wants to thank Joshua Bowman, Vincent Delecroix, Rodrigo Treviño, Chenxi Wu, Lucien Clavier, and Rodolfo Rios-Zertuche for helpful discussions and hints.

1 Basics on translation surfaces and their singularities

As there are several gradations in generalizing the notion of finite translation surfaces we first make clear the definition that we use in this article.

Definition 1 ((Finite) Translation surface)
A *translation surface* (X, ω) is a connected surface X with a translation structure ω on X, i.e. a maximal atlas on X so that the transition functions are translations. Via the translation structure we can pull back the Euclidean metric from $\mathbb{R}^2$ to X.

The translation surface (X, ω) is called *finite* if the metric completion $\overline{X}$ is a compact surface.

Note that for a translation surface which is not finite there exist examples where the metric completion is not compact as well as where it is not a surface. For the first one consider the Eudlidean plane $\mathbb{R}^2$ and for the second one the baker's map surface studied by Chamanara in [Cha04] or all of the examples in section 4.

For a translation surface (X, ω) we call the points in $\overline{X} \setminus X$ *singularities* of (X, ω). In contrast to finite translation surfaces very exciting behaviour of singularities can occur. We will distinguish between different types of singularities:

Definition 2 (Cone angle, infinite angle, and wild singularities)
Let (X, ω) be a translation surface and σ a singularity of (X, ω).

(i) The singularity σ is called *cone angle singularity of multiplicity* $k > 0$ if there exist

- $\epsilon > 0$,
- an open neighborhood B of σ in $\overline{X}$, and
- a k-cyclic translation covering from $B \setminus \{\sigma\}$ to the punctured disk $B(0, \epsilon) \setminus \{0\} \subseteq \mathbb{R}^2$.

If $k = 1$, σ is also called *flat point* or *removable singularity.*

(ii) The singularity σ is called *infinite angle singularity* or *cone angle singularity of multiplicity* ∞ if there exist

- $\epsilon > 0$,
- an open neighborhood B of σ in $\overline{X}$, and
- an infinite cyclic translation covering from $B \setminus \{\sigma\}$ to the punctured disk $B(0, \epsilon) \setminus \{0\} \subseteq \mathbb{R}^2$.

(iii) The singularity σ is called *wild* if it is neither a cone angle nor an infinite angle singularity.

For cone angle and infinite angle singularities, all topological information is decoded in the multiplicity. To understand wild singularities in a similar way we have to describe them more detailed. The first attempt of that was done by Bowman and Valdez in [BV13]. We will recall it very briefly:

Definition 3 (Space of linear approaches)
Let (X, ω) be a translation surface, $x \in \overline{X}$, and $\epsilon > 0$. We define

$$\mathcal{L}^\epsilon(x) \coloneqq \left\{ \gamma \colon (0, \epsilon) \to X : \gamma \text{ is a geodesic curve and } \lim_{t \to 0} \gamma(t) = x \right\}.$$

As we can deduce directly from Definition 2 there is no good choice of a global ϵ if x is a wild singularity so we consider equivalence classes instead of curves: $\gamma_1 \in \mathcal{L}^\epsilon(x)$ and $\gamma_2 \in \mathcal{L}^{\epsilon'}(x)$ are called *equivalent* if $\gamma_1(t) = \gamma_2(t)$ for all $t \in (0, \min\{\epsilon, \epsilon'\})$.

The space

$$\mathcal{L}(x) \coloneqq \bigsqcup_{\epsilon > 0} \mathcal{L}^\epsilon(x) / \sim$$

is called *space of linear approaches of* x and the equivalence class $[\gamma]$ of $\gamma \in \mathcal{L}^\epsilon(x)$ is called *linear approach* to the point x.

Definition 4 (Topology on $\mathcal{L}(x)$)
For $x \in \overline{X}$ and $\epsilon > 0$ we can define the *uniform metric* on $\mathcal{L}^\epsilon(x)$ using the translation metric d_X on X:

$$d_\epsilon(\gamma_1, \gamma_2) = \sup_{0 < t < \epsilon} d_X\left(\gamma_1(t), \gamma_2(t)\right)$$

The metric also defines a topology on $\mathcal{L}^\epsilon(x)$. Every $\mathcal{L}^\epsilon(x)$ can be embedded in $\mathcal{L}(x)$ and also in $\mathcal{L}^{\epsilon'}(x)$ for all $\epsilon' > 0$ with $\epsilon \geq \epsilon'$. By this we obtain a direct system, i.e. the composition of two such embeddings $\mathcal{L}^\epsilon(x) \hookrightarrow \mathcal{L}^{\epsilon'}(x)$, $\mathcal{L}^{\epsilon'}(x) \hookrightarrow \mathcal{L}^{\epsilon''}(x)$ is equal to the embedding $\mathcal{L}^\epsilon(x) \hookrightarrow \mathcal{L}^{\epsilon''}(x)$. So we can define the *final topology* on $\mathcal{L}(x)$. This is the finest topology so that all embeddings $\mathcal{L}^\epsilon(x) \hookrightarrow \mathcal{L}(x)$ are continuous.

For a cone angle and infinite angle singularity, we can "go around" this singularity to find out the multiplicity. For a wild singularity, there is no possibility to "go around" in a fixed distance and no meaningful way to define the multiplicity. So, we want to define an object that is feasible for all type of points in $\overline{X}$.

First, we consider the infinite strip $\{z \in \mathbb{C} : \mathrm{Re}(z) < \log \epsilon, \mathrm{Im}(z) \in I\} \subseteq \mathbb{C}$ for a fixed ϵ and a fixed (generalized) interval I.

Via the injective map

$$f \colon \mathbb{C} \to (\mathbb{C} \setminus \{0\}) \times \mathbb{R}, z \mapsto (e^z, \mathrm{Im}(z))$$

we can spiral the strip around the missing point 0. The image U of the strip under f is endowed with the pullback of the Euclidean metric on $\mathbb{C}$ via the canonical map $U \to B(0, \epsilon)$. Some examples are sketched in Figure 1.

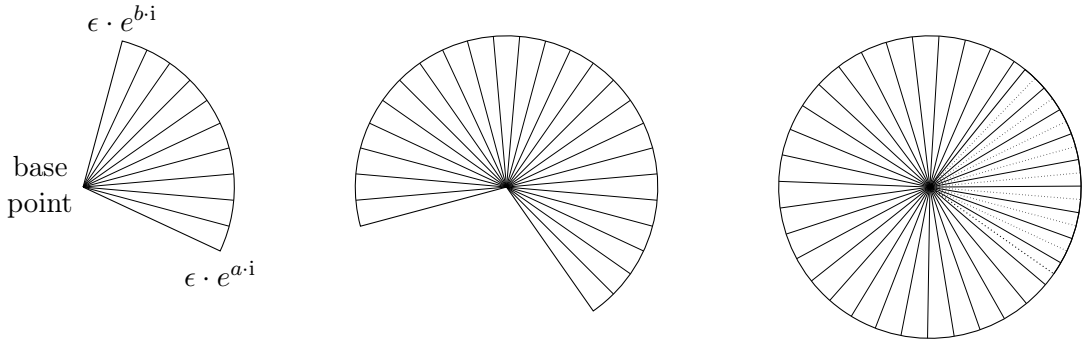

Figure 1: Examples for U resp. $i_\epsilon(U)$ with $I = [a, b]$.

Definition 5 (Angular sector)
An *angular sector* is a tripel $(I, \epsilon, i_\epsilon)$ of a generalized interval I (i.e. a non-empty connected subset of $\mathbb{R}$), $\epsilon > 0$, and an isometric embedding i_ϵ of

$$U \coloneqq f\left(\{z \in \mathbb{C} : \mathrm{Re}(z) < \log \epsilon, \mathrm{Im}(z) \in I\}\right)$$

in X.

For an angular sector $(I, \epsilon, i_\epsilon)$ and $y \in I$,

$$\lim_{x \to -\infty} (i_\epsilon \circ f)(x + \mathrm{i}y)$$

is a point in $\overline{X}$ and independent of y. This point is called *base point* of the angular sector $(I, \epsilon, i_\epsilon)$.

In the next definition, we use angular sectors to generalize the notion of "going around a singularity". The formal tool for this is an equivalence relation on the space of linear approaches of a point. In a sloppy way, we can describe it as two linear approaches being "contained" in the image of a given angular sector.

Definition 6 (Rotational component)
Let $x \in \overline{X}$ and $[\gamma_1], [\gamma_2] \in \mathcal{L}(x)$. The linear approaches $[\gamma_1]$ and $[\gamma_2]$ are called *equivalent* if there exist

- $\epsilon > 0$,
- representants $\gamma_1, \gamma_2 \in \mathcal{L}^\epsilon(x)$,
- an angular sector $(I, \epsilon, i_\epsilon)$ with base point x,
- and $y_1, y_2 \in I$

so that we have

$$\mathrm{im}(\gamma_j) = (i_\epsilon \circ f)\left(\{z \in \mathbb{C} : \mathrm{Re}(z) < \log \epsilon, \mathrm{Im}(z) = y_j\}\right) \text{ for } j = 1, 2.$$

The equivalence class $\overline{[\gamma]}$ of $[\gamma] \in \mathcal{L}(x)$ is called *rotational component* of x.

For an angular sector $(I, \epsilon, i_\epsilon)$ with base point $x \in \overline{X}$, we can assign a linear approach $[\gamma_y]$ to each $y \in I$ similar to Definition 6: the set

$$(i_\epsilon \circ f)\left(\{z \in \mathbb{C} : \mathrm{Re}(z) < \log \epsilon, \mathrm{Im}(z) = y\}\right) \subseteq X$$

defines a representant γ_y of $[\gamma_y]$.

The set $V(I, \epsilon, i_\epsilon) \coloneqq \{[\gamma_y] : y \in I\}$ of all linear approaches assigned to $(I, \epsilon, i_\epsilon)$ is a subset of the space of linear approaches $\mathcal{L}(x)$. All such linear approaches are contained in the same rotational component $\overline{[\gamma]}$ where $\overline{[\gamma]} \coloneqq \overline{[\gamma_y]}$ for any $y \in I$. Moreover, the map

$$\varphi_{V(I,\epsilon,i_\epsilon)} \colon V(I, \epsilon, i_\epsilon) \to \mathbb{R},\ [\gamma_y] \mapsto y$$

is one-to-one from $V(I, \epsilon, i_\epsilon)$ to I.

So we can define a topology on the rotational component $\overline{[\gamma]}$ by considering all angular sectors $(I, \epsilon, i_\epsilon)$ with open interval I and base point x. We choose the set of the corresponding $V(I, \epsilon, i_\epsilon)$ as a basis of the topology. Then any such $\varphi_{V(I,\epsilon,i_\epsilon)}$ is a homeomorphism and

$$\{(V(I, \epsilon, i_\epsilon), \varphi_{V(I,\epsilon,i_\epsilon)}) : (I, \epsilon, i_\epsilon) \text{ an angular sector with base point } x, I \text{ open}\}$$

forms an atlas of $\overline{[\gamma]}$. Therefore, $\overline{[\gamma]}$ is a one-dimensional manifold, more precisely: it has a one-dimensional translation structure.

If $\overline{[\gamma]}$ contains boundary points we additionally consider sets $V(I, \epsilon, i_\epsilon)$ with a half-closed interval I as open sets if a boundary point of $\overline{[\gamma]}$ is contained in $V(I, \epsilon, i_\epsilon)$ and is mapped to the boundary point of I by $\varphi_{V(I,\epsilon,i_\epsilon)}$. In this situation, $\overline{[\gamma]}$ is a one-dimensional manifold with boundary.

Using the translation structure of $\overline{[\gamma]}$, we can pull back the Euclidean metric from $\mathbb{R}$ to $\overline{[\gamma]}$ and for this we can measure the *length of a rotational component*. For instance, the length of the only rotational component of a cone angle singularity of multiplicity k is $2\pi k$.

2 Criterion for infinite genus

As we want to show that a surface has infinite genus we need a feasible criterion for the infinity of genus. For this we start with a well-known criterion and specialize it for the case of translation surfaces.

Definition 7 (Non-separating curves)
Let X be a connected surface.

(i) A simple closed curve $\gamma\colon [0,l] \to X$ is called *non-separating* if X with the image of γ removed is connected.

(ii) The set of simple closed curves $\{\gamma_1, \dots, \gamma_n\}$ is called *non-separating* if X with the images of $\gamma_1, \dots, \gamma_n$ removed is connected.

Similarly, we say that a closed curve in $\overline{X}$ which connects a singularity to itself is *non-separating* if $\overline{X} \setminus \operatorname{im}(\gamma)$ is connected.

Proposition 8 (Genus defines maximum number of non-separating curves)
Let X be a connected surface of genus g. Then the following holds:

(i) The maximum cardinality of a non-separating set of disjoint curves in X is g.

(ii) The maximum cardinality of a non-separating set of curves in X is $2g$.

Proof. Sometimes, (i) is taken as a definition of genus. A proof of (i) and (ii) is given in [Ful97, Chapter 17] using fundamental polygons. □

We now want to prove a criterion for curves to be non-separating. For that we use curves that connect one side of a curve to the other side. We call these curves *left-to-right curves*. This will be made precise in the following definition:

Definition 9 (Left-to-right curves)
Let X be a connected surface and $\gamma, \gamma_1, \dots, \gamma_n$ simple closed curves in X.

(i) Let $\epsilon > 0$ be small enough so that the ϵ-neighborhood N of γ is a tubular neighborhood. Then N is topologically an annulus. So $N \setminus \operatorname{im}(\gamma)$ consists of two connected components N_l and N_r. Considering the underlying orientation of $\operatorname{im}(\gamma)$ we will call points in N_l and N_r *points on the left of* γ and *points on the right of* γ, respectively.

(ii) A curve in $X \setminus \operatorname{im}(\gamma)$ from a point on the left of γ to a point on the right of γ is called *left-to-right curve* of γ.

(iii) Suppose that $\{\gamma_1, \dots, \gamma_{n-1}\}$ is non-separating and that furthermore the curves intersect pairwise in exactly one point P. Furthermore, suppose that γ_n intersects $\gamma_1, \dots, \gamma_{n-1}$ exactly in P.

Choose a tubular neighborhood N of γ_n and let N_l and N_r be as before. We have that $N_l \setminus (\operatorname{im}(\gamma_1) \cup \dots \cup \operatorname{im}(\gamma_{n-1}))$ consists of one or more connected components (see Figure 2). The boundary of such a connected component consists of a subset of the boundary of N and of subsets of the images of some γ_i. As the curves intersect in exactly one point, there is only one connected component N_l^1 whose boundary additionally contains $\operatorname{im}(\gamma_n)$.

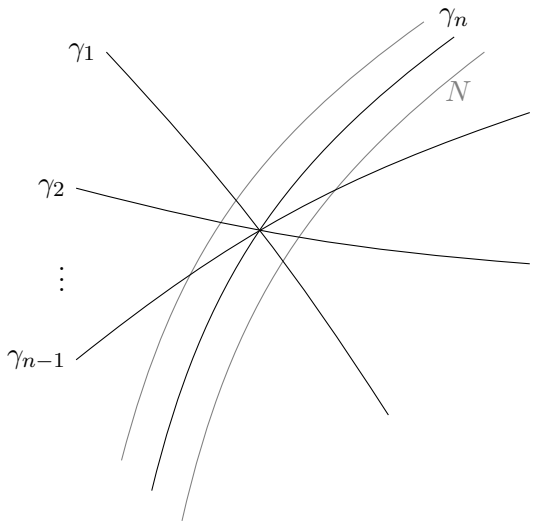

Figure 2: The connected components N_l^1 and N_r^1 are the ones of which only parts are drawn.

We call a point in this connected component *point on the left of γ_n with respect to $\gamma_1, \ldots, \gamma_{n-1}$*. Similarly we can define *points on the right of γ_n with respect to $\gamma_1, \ldots, \gamma_{n-1}$*. Then a curve in $X \setminus (\mathrm{im}(\gamma_1) \cap \ldots \cap \mathrm{im}(\gamma_n))$ is called *left-to-right curve of γ_n with respect to $\gamma_1, \ldots, \gamma_{n-1}$* if it connects a point on the left of γ_n with respect to $\gamma_1, \ldots, \gamma_{n-1}$ to a point on the right of γ_n with respect to $\gamma_1, \ldots, \gamma_{n-1}$.

(iv) We say *a set of curves $\{\gamma_1, \ldots, \gamma_n\}$ has left-to-right curves* if every curve has a left-to-right curve with respect to the other ones.

(v) Now let (X, ω) be a translation surface with discrete singularities. For a curve in $\overline{X}$ like a saddle connection we cannot find such a tubular neighborhood N but use a slightly different neighborhood instead. So we can define left-to-right curves for a special type of curves in $\overline{X}$ in a similar way while avoiding the singularity:

Let σ be a singularity, γ be a curve in $X \cup \{\sigma\}$ whose image contains σ exactly as start and end point, and l be the length of γ. Furthermore, let $\epsilon > 0$ be small enough and consider the set N which is the union of $B(\sigma, \epsilon)$ and the open ϵ-neighborhood $\widetilde{N}$ of the segment $\gamma([\epsilon, l - \epsilon])$. In this situation, we call points in $\widetilde{N} \setminus \mathrm{im}(\gamma)$ *points on the left of γ* and *points on the right of γ*, with respect to the orientation. Then a curve δ in $X \setminus \mathrm{im}(\gamma)$ is called a *left-to-right curve* of γ if it connects a point on the left of γ to a point on the right of γ.

In analogy to the situation of simple closed curves, we can define *left-to-right curves with respect to a set* of this type of curves in $\overline{X}$ where the curves intersect exactly in the singularity. Furthermore, we can define what it means for a set of these curves to have left-to-right curves.

Note that the existence of left-to-right curves does not depend on the ϵ (as long as ϵ is small) or on the tubular neighborhood that we choose.

With these left-to-right curves we can now formulate a criterion for a curve and a set of curves to be non-separating.

Lemma 10 (Criterion for simple closed curves to be non-separating). Let X be a connected surface, $n \geq 2$, and $\gamma, \gamma_1, \ldots, \gamma_n$ simple closed curves in X.

(i) The curve γ is non-separating if and only if it has a left-to-right curve.

(ii) The set $\{\gamma_1, \ldots, \gamma_n\}$ is non-separating if and only if the set has left-to-right curves.

Proof. (i) Although this statement is a special case of the statement about non-separating sets we will carry out the proof. In the next proof we will use it as base case for an induction argument.

If γ is non-separating we can take any point on the right and any point on the left of γ and as $X \setminus \mathrm{im}(\gamma)$ is connected there exists a curve connecting these two points without intersecting γ.

Now assume we have such a left-to-right curve δ connecting x_l and x_r. Choose two points $x_1, x_2 \in X \setminus \mathrm{im}(\gamma)$. We have to show that there exists a curve β in $X \setminus \mathrm{im}(\gamma)$ that connects x_1 and x_2.

Let N, N_l and N_r be as in Definition 9. As X is connected there exists a curve β' in X that connects x_1 to x_2. If β' is disjoint from N we can choose $\beta := \beta'$. If not let β'_+ be the subcurve of β' from x_1 to the first intersection of β' and ∂N and let β'_- be the subcurve of β' from the last intersection of β' and ∂N to x_2 (see Figure 3 for a sketch).

Case 1: The endpoint of β'_+ and the startpoint of β'_- belong both to ∂N_r or both to ∂N_l. Then we can choose a curve between these two points in the connected set N_r or N_l, respectively, and the concatenation of this curve with β'_+ and β'_- gives us a curve β as desired.

Case 2: The endpoint of β'_+ belongs to ∂N_l or to ∂N_r and the startpoint of β'_- belongs to the other. Without loss of generality, let β'_+ end in ∂N_l and β'_- start in ∂N_r. Again, as N_l and N_r are connected we find curves connecting the endpoint of β'_+ to x_l and connecting x_r to the startpoint of β'_-. Concatenating all these curves with the left-to-right curve δ in the correct order gives us a curve β as desired.

(ii) If $\{\gamma_1, \ldots, \gamma_n\}$ is a non-separating set of curves we can take any point on the right and any point on the left of γ_i with respect to $\gamma_1, \ldots, \gamma_{i-1}, \gamma_{i+1}, \ldots, \gamma_n$ and as $X \setminus \{\mathrm{im}(\gamma_1) \cup \ldots \cup \mathrm{im}(\gamma_n)\}$ is connected there exists a curve connecting these two points without intersecting one of the given curves.

Now assume the set $\{\gamma_1, \ldots, \gamma_n\}$ has left-to-right curves. We will use induction on n to show that the set is non-separating. The base case $n = 1$ is done in (i) so assume that for $n \geq 2$ the set $\{\gamma_1, \ldots, \gamma_{n-1}\}$ is non-separating and γ_n has a

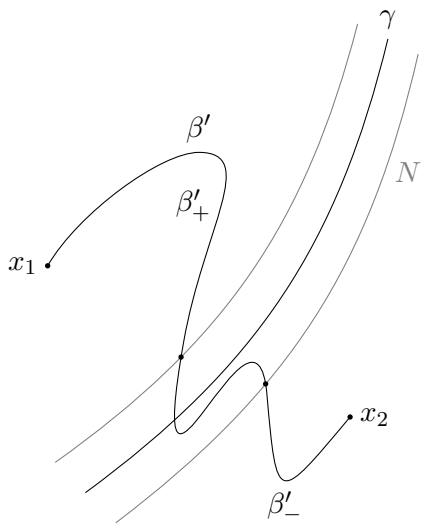

Figure 3: This configuration of γ and β' is treated in case 2.

left-to-right curve δ_n with respect to $\gamma_1, \dots, \gamma_{n-1}$ connecting a point x_l on the left of γ_n to a point x_r on the right of γ_n.

Again, we have to show that for two points $x_1, x_2 \in X \setminus (\operatorname{im}(\gamma_1) \cup \dots \cup \operatorname{im}(\gamma_n))$ there exists a curve β_n in $X \setminus (\operatorname{im}(\gamma_1) \cup \dots \cup \operatorname{im}(\gamma_n))$ that connects x_1 and x_2. By the induction hypothesis there exists a curve β'_n in $X \setminus (\operatorname{im}(\gamma_1) \cup \dots \cup \operatorname{im}(\gamma_{n-1}))$ that connects x_1 and x_2.

If β'_n does not intersect γ_n we can choose $\beta_n \coloneqq \beta'_n$. If it does, it also has to intersect one of the connected components N_l^1 or N_r^1 of $N \setminus (\operatorname{im}(\gamma_1) \cup \dots \cup \operatorname{im}(\gamma_n))$. This is because the boundaries of all other connected components contain only one point of $\operatorname{im}(\gamma_n)$ which is the intersection point. Now let β'_+ be the subcurve of β'_n from x_1 to the first intersection of β'_n and $\partial(N_l^1 \cup N_r^1)$ and let β'_- be the subcurve of β'_n from the last intersection of β'_n and $\partial(N_l^1 \cup N_r^1)$ to x_2. Now we can proceed as in (i) and can construct a curve β_n in $X \setminus (\operatorname{im}(\gamma_1) \cup \dots \cup \operatorname{im}(\gamma_n))$ that connects x_1 and x_2. □

Remark 11 (Lemma 10 for closed curves in $\overline{X}$). It follows by the same arguments, that the previous criterion is also true for saddle-connections or, more generally, closed curves in $\overline{X}$ whose image contains singularities only as start and end point. While the beginning of the proof of statement (i) is completely the same, we have to choose an ϵ small enough so that all intersection points of β' and γ are in the ϵ-neighborhood $\widetilde{N}$ of the segment $\gamma([\epsilon, l - \epsilon])$. Then we can use $\widetilde{N}$ instead of N and finish the proof in the same way as before. The replacement of N by $\widetilde{N}$ also makes the proof of statement (ii) work for this

type of curves as the components of $\widetilde{N} \setminus \mathrm{im}(\gamma)$ can also play the role of the connected components N_l^1 and N_r^1.

As "genus" is a concept for surfaces we have to consider curves in X instead of curves in $\overline{X}$ to determine infinite genus. Because of this, we show how to replace one kind of curves by other curves without disturbing the left-to-right curves in the next lemma. So we have to make sure the existence of short curves between given points which we will write down as a proposition for the sake of references.

Proposition 12 (Balls around singularities are path-connected)
Let (X, ω) be a translation surface, σ a singularity, and $\epsilon > 0$. Then the ϵ-neighborhood $B(\sigma, \epsilon)$ of σ is path-connected.

Proof. Let $x, y \in B(\sigma, \epsilon)$ and $\delta > 0$ such that $d(x, \sigma) < \epsilon - 3\delta$ and $d(y, \sigma) < \epsilon - 3\delta$. Furthermore let $(z_n)_{n\in\mathbb{N}} \subseteq X$ be a Cauchy sequence converging to σ and $N \in \mathbb{N}$ with $d(\sigma, z_N) < \delta$. Then we have $d(x, z_N) \leq d(x, \sigma) + d(\sigma, z_N) < \epsilon - 2\delta$ and $d(y, z_N) < \epsilon - 2\delta$.

By the definition of the distance in X via the infimum of curve lengths we know that there exists a curve γ_x in X which connects x to z_N and has at most length $\epsilon - \delta$. This means that for every point in the image of γ_x the distance to σ is at most $d(\sigma, z_N) + (\epsilon - \delta) = \epsilon$. Hence, γ_x is a curve in $B(\sigma, \epsilon)$.

In the same way we can define γ_y as a curve in $B(\sigma, \epsilon)$ from y to z_N and by concatenation of γ_x and the reversed curve of γ_y we have a curve from x to y in $B(\sigma, \epsilon)$. This means, $B(\sigma, \epsilon)$ is path-connected. □

Lemma 13 (Non-separating curves in $\overline{X}$ give rise to non-separating curves in X). Let (X, ω) be a translation surface, σ a singularity, and $\gamma_1, \ldots, \gamma_n$ a non-separating set of curves in $X \cup \{\sigma\}$ whose images contain σ exactly as start and end points.

Then there also exists a set of curves $\gamma_1', \ldots, \gamma_n'$ in X that is non-separating.

Proof. By assumption there exist left-to-right curves δ_i of γ_i with respect to $\gamma_1, \ldots, \gamma_{i-1}$, $\gamma_{i+1}, \ldots, \gamma_n$ for all $i \in \{1, \ldots, n\}$. Choose ϵ small enough so that the ϵ-neighborhood $B(\sigma, \epsilon)$ of σ avoids all δ_i and so that $\partial B(\sigma, \epsilon)$ intersects $\mathrm{im}(\gamma_i)$ at least two times for all $i \in \{1, \ldots, n\}$. For every γ_i, the first intersection point of $\partial B(\sigma, \epsilon)$ and $\mathrm{im}(\gamma_i)$ (with respect to the orientation of $\mathrm{im}(\gamma_i)$) is called x_i^+ and the last intersection point is called x_i^-.

We will now replace the curves $\gamma_1, \ldots, \gamma_n$ by curves in X that have similar properties. For this choose a point $x \in B(\sigma, \epsilon) \setminus \{\sigma\}$ that will play the role of the current intersection point σ. Because of the path connectedness of $B(\sigma, \epsilon)$ we have a curve in $B(\sigma, \epsilon) \setminus \{\sigma\}$ between x and x_1^+ and a curve between x_1^- and x. Now let γ_1' be the simple closed curve that consists of the curve between x and x_1^+, the subcurve of γ_1 from x_1^+ to x_1^-, and the curve between x_1^- and x. If γ_1' intersects itself we can smooth the crossing by joining other pairs of subcurves at the crossing. So without loss of generality γ_1' is simple. Also, as ϵ is chosen small enough, the curve δ_1 is still a left-to-right curve of γ_1' with respect to $\gamma_2, \ldots, \gamma_n$.

Now we do the same construction for the rest of the curves successively: For the construction of γ_i' we need a curve from x_i^+ to the prospective intersection point x that

does not leave $B(\sigma, \epsilon)$ and does not intersect the curves $\gamma'_1, \dots, \gamma'_{i-1}$. We can find it by taking any curve from x_i^+ to x in $B(\sigma, \epsilon) \setminus \{\sigma\}$ and instead of crossing some γ'_j we travel next to it until we reach x. Then we define γ'_i as the closed curve that consists of the subcurve of γ_i from x_i^+ to x_i^- and a curve between x_i^+ and x as described and similarly a curve between x_i^- and x. Again, the curve δ_i is still a left-to-right curve of γ'_i with respect to $\gamma_1, \dots, \gamma_{i-1}, \gamma_{i+1}, \dots, \gamma_n$.

So we have a set of simple closed curves $\gamma'_1, \dots, \gamma'_n$ that are intersecting exactly in x. Also, the set has left-to-right curves $\delta_1, \dots, \delta_n$ so it is non-separating by Lemma 10. □

By the criterion in Lemma 10, we can show that a set of curves is non-separating but so far we don't have candidates of non-separating curves where we could use the criterion. Therefore we introduce a generalization of the well-known fact that saddle connections of finite translation surfaces are non-separating:

Lemma 14 (Saddle connections are non-separating). Let (X, ω) be a translation surface so that for two directions θ the geodesic flow F_θ is recurrent for almost every point.

Let γ be a saddle connection starting and ending at the same singularity. Then γ is non-separating.

Proof. Consider a geodesic segment $s \subseteq \mathrm{im}(\gamma)$ and the geodesic flow F_θ in a direction θ that is transversal to the direction of γ and so that the flow is recurrent for almost every point. So there exists a point $p \in s$ that returns to s under the flow F_θ after time t_1. In particular, there exists a time t_0 with $0 < t_0 \leq t_1$ such that $F_\theta(p, t_0) \in \mathrm{im}(\gamma)$ for the first time. Additionally, as γ is geodesic and F_θ is a geodesic flow the curve $\delta\colon [0, t_0] \to X$, $t \mapsto F_\theta(p, t)$ is arriving at $\mathrm{im}(\gamma)$ from the other side than it is leaving. Then there exists an $\epsilon > 0$ so that the curve $\delta_\epsilon\colon [\epsilon, t_0 - \epsilon] \to X$, $t \mapsto F_\theta(p, t)$ is a curve in $X \setminus \mathrm{im}(\gamma)$. The curve δ_ϵ or its reversed curve connects a point on the left side of γ to a point on the right side of γ, so it is a left-to-right curve of γ. Hence, γ is non-separating. □

In the next section we will construct a set of saddle connections carefully. To show that the existence of this set is enough for our purpose we can use the following summary of the section:

Proposition 15 (Saddle connections and infinite genus)
Let (X, ω) be a translation surface. Suppose for every $n \in \mathbb{N}$ there exists a set of n saddle connections from a singularity to itself so that the saddle connections intersect exactly in the singularity and the set has left-to-right curves. Then X has infinite genus.

Proof. As the saddle connections are curves as in the conditions of Lemma 13 we have a non-separating set of n curves in X for every $n \in \mathbb{N}$. According to Proposition 8 this means that the genus of X exceeds every number n so X has infinite genus. □

3 Short saddle connections

In the next definition we will have to use that saddle connections can be seen as linear approaches. In the following remark this is made explicit:

Remark 16 (Saddle connections and linear approaches). Recall that every saddle connection $\gamma\colon [0,l] \to X$ is geodesic and has by definition an orientiation. So the curve $\gamma_{|(0,l)}$ and its reversed curve define two linear approaches, one belonging to emanating from a singularity and one belonging to going into a singularity. For a saddle connection γ we call the first linear approach $[\gamma_+]$ and the second one $[\gamma_-]$.

Similar to the well-known injectivity radius, we will now define the immersion radius of a point resp. a curve. In contrast to the injectivity radius, we allow that the image of the disk we map into X overlaps itself.

Definition 17 (Immersion radius)
Let (X,ω) be a translation surface with discrete singularities and σ a singularity.

(i) For a point $x \in X$ we define the *immersion radius* $\mathrm{ir}(x)$ as

$$\mathrm{ir}(x) \coloneqq d(x, \overline{X} \setminus X) \in (0,\infty].$$

This is well defined as $\overline{X}$ is a metric space. Note that the open $\mathrm{ir}(x)$-neighborhood of x is locally flat, i.e. it does not contain a singularity.

(ii) For a curve $\gamma\colon [0,l] \to X$ we define the *immersion radius* $\mathrm{ir}(\gamma)$ by

$$\mathrm{ir}(\gamma) \coloneqq \inf\{\mathrm{ir}(\gamma(t)) : t \in [0,l]\}.$$

As the image of γ is compact we can cover it by finitely many balls $B(\gamma(t_i), \mathrm{ir}(\gamma(t_i)))$ around points $\gamma(t_i)$ on the curve. Then the distance of $\partial \bigcup B(\gamma(t_i), \mathrm{ir}(\gamma(t_i)))$ to $\mathrm{im}(\gamma)$ is the minimum of finitely many positive numbers and so it is positive, too. As this distance is a lower bound for the immersion radius of the curve this means that $\mathrm{ir}(\gamma) > 0$ still holds and again the open $\mathrm{ir}(\gamma)$-neighborhood of $\mathrm{im}(\gamma)$ is locally flat. In particular, every saddle connection that intersects γ has at least length $\mathrm{ir}(\gamma)$.

(iii) A saddle connection has to be seen either as a curve in $\overline{X}$ or as an open geodesic. Because of that we have to slightly modify the notion of immersion radius. As this modification will only work for saddle connections for which both corresponding linear approaches (as in Remark 16) belong to a rotational component whose length is greater than π we will restrict on this type of saddle connections:

Let $\gamma\colon [0,l] \to \overline{X}$ be a saddle connection and let $[\gamma_+]$ be as in Remark 16. By our condition on the rotational component there exists an angular sector $(I_+, 2\epsilon_+, i_{2\epsilon_+})$ so that the length of I is greater than π and $[\gamma_+]$ is a linear approach assigned to $(I_+, 2\epsilon_+, i_{2\epsilon_+})$ as in the remark after Definition 6. Let γ_+ be the representant of $[\gamma_+]$ in $\mathcal{L}^{\epsilon_+}(\sigma)$. Because the length of I is at least π, the image of γ_+ is contained in the union of a locally flat, open halfdisk and σ (see Figure 4). Suppose there exists a saddle connection that intersects γ_+. Then it has to start outside of the image of $i_{2\epsilon_+}$ and go through the half-annulus-like set

$$\left(i_{2\epsilon_+} \circ f\right)\left(\{z \in \mathbb{C} : \log \epsilon_+ < \mathrm{Re}(z) < \log 2\epsilon_+, \mathrm{Im}(z) \in I_+\}\right)$$

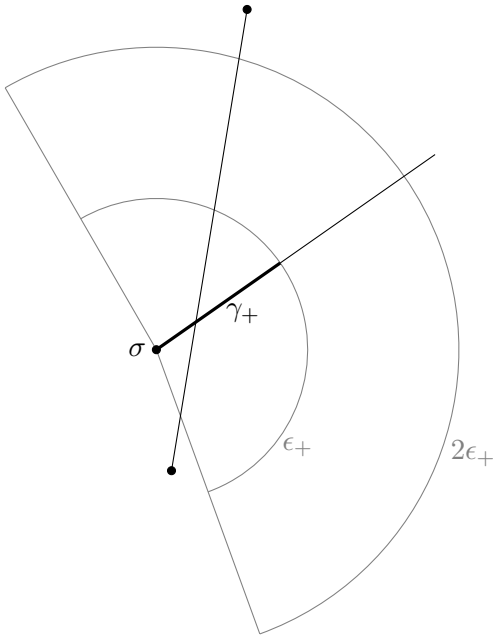

Figure 4: Any saddle connection that intersects the given saddle connection γ has to have a certain length.

to intersect the image of γ_+. This means that the intersecting saddle connection has at least length ϵ_+.

Define γ_- and ϵ_- similarly by using the linear approach $[\gamma_-]$ as in Remark 16. Let $\gamma_c \coloneqq \gamma_{|[\epsilon_+, l-\epsilon_-]}$. This is a curve as in (ii) with a compact image and a well-defined immersion radius $\epsilon_c \coloneqq \mathrm{ir}(\gamma_c) > 0$.

Then $\min\{\epsilon_+, \epsilon_c, \epsilon_-\}$ is a lower bound for the length of saddle connections that intersect γ. So we take this as a generalization of immersion radii of curves to the set of saddle connections. For a saddle connection γ contained in rotational components of length greater than π, we define the *(generalized) immersion radius* $\mathrm{ir}(\gamma)$ by

$$\mathrm{ir}(\gamma) \coloneqq \sup\{\min\{\epsilon_+, \epsilon_c, \epsilon_-\} : \epsilon_+, \epsilon_- \text{ small enough so that angular sectors as described before exist}\}.$$

Remark 18 (Justification of the term "immersion radius"). In the case of a point x, the term "immersion radius" is certainly justified as a disk of radius $\mathrm{ir}(x)$ can be immersed and the image under the immersion is a locally flat neighborhood of the point x. Similarly, for a closed curve a cylinder of height $2 \cdot \mathrm{ir}(\gamma)$ and circumference the length of γ can be immersed and the image will be a locally flat tubular neighborhood of the curve.

In the case of saddle connections, we can immerse a trapezoid of height $2 \cdot \mathrm{ir}(\gamma)$ so that the median has the same length as the saddle connection (see Figure 5). The image will

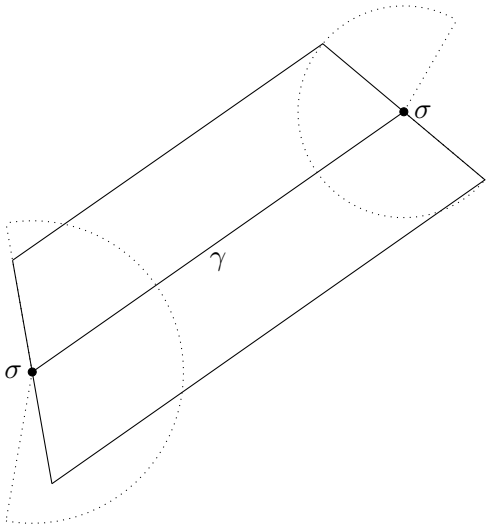

Figure 5: An example of a trapezoidal neighborhood of a saddle connection γ.

form a neighborhood of the interior of the saddle connection and around the singularities we have the images of angular sectors of length π.

Lemma 19 (Immersion radius is Lipschitz continuous). The map $\mathrm{ir}\colon X \to (0,\infty], x \mapsto \mathrm{ir}(x)$ is Lipschitz continuous with constant 1. In particular, the map is continuous.

Proof. For every $x_1, x_2 \in X$ and $\sigma \in \overline{X} \setminus X$ we have

$$\mathrm{ir}(x_1) = d(x_1, \overline{X} \setminus X) \leq d(x_1, \sigma) \leq d(x_1, x_2) + d(x_2, \sigma).$$

Therefore it follows

$$\mathrm{ir}(x_1) \leq d(x_1, x_2) + \inf_{\sigma \in \overline{X} \setminus X} d(x_2, \sigma) = d(x_1, x_2) + \mathrm{ir}(x_2)$$

and interchanging x_1 and x_2 gives us $|\mathrm{ir}(x_1) - \mathrm{ir}(x_2)| \leq d(x_1, x_2)$. So the map is Lipschitz continuous with Lipschitz constant 1. □

Definition 20 (Immersion radius along a linear approach)
Let (X, ω) be a translation surface with discrete singularities, σ a singularity and $\gamma\colon (0,\epsilon) \to X$ a representant of the linear approach $[\gamma] \in \mathcal{L}(\sigma)$.

For every $t \in (0,\epsilon)$ the immersion radius of $\gamma(t)$ is greater 0 but at most $d(\gamma(t), \sigma) \leq t$. So we can define the map

$$\mathrm{ir}_\gamma\colon (0,\epsilon) \to (0,\epsilon),\ t \mapsto \mathrm{ir}(\gamma(t)).$$

We will use the immersion radius soon as a tool to find saddle connections which do not intersect a given set of curves. However, first we give a weaker statement where we don't need this tool.

Proposition 21 (Weak existence of short saddle connections)
Let (X, ω) be a translation surface with discrete singularities and let σ be a wild singularity of (X, ω). Then for every $\epsilon > 0$ there exists a saddle connection connecting σ to itself of length less than ϵ.

Proof. As the singularities are discrete there exists a $\delta > 0$ so that σ is the only singularity in $B(\sigma, \delta)$. We now fix an $\epsilon > 0$ with $\epsilon < \delta$.

If for every linear approach $[\gamma]$ to σ there existed a representant $\gamma \in \mathcal{L}^\epsilon(\sigma)$, then we would have a cyclic translation covering from $B(\sigma, \epsilon) \setminus \{\sigma\}$ to the punctured disk $B(0, \epsilon) \setminus \{0\} \subseteq \mathbb{R}^2$. This means that σ wouldn't be a wild singularity.

So there exists a linear approach $[\gamma]$ to σ for which the longest representant γ has a length less than ϵ. This means the closure of the image of γ in $\overline{X}$ is a saddle connection from σ to σ as desired. □

We can formulate a stronger version of the previous statement:

Proposition 21' (Strong existence of short saddle connections)
Let (X, ω) be a translation surface with discrete singularities and let σ be a wild singularity of (X, ω) such that none rotational component of σ has length π. Then for every $\epsilon > 0$ there exists a saddle connection connecting σ to itself of length less than ϵ and so that both corresponding linear approaches are contained in a rotational component of length greater than π.

Proof. Again, we fix $\delta > 0$ so that σ is the only singularity in $B(\sigma, \delta)$ and $\epsilon > 0$ with $\epsilon < \delta$. To fulfill the part of the statement about the linear approaches we have to choose the saddle connection a bit more carefully as before.

Similar to the previous proof, there exists a linear approach $[\gamma]$ for which the longest representant γ has a length $l < \frac{\epsilon}{2}$. For a time $t \in (0, l)$ with $\mathrm{ir}_\gamma(t) < t$ there exists a geodesic in $\overline{X}$ of length $\mathrm{ir}_\gamma(t)$ connecting $\gamma(t)$ to a singularity. By $t + \mathrm{ir}_\gamma(t) < \frac{\epsilon}{2} + \frac{\epsilon}{2} < \delta$, this singularity is again σ (see Figure 6).

To prove the statement we will show the existence of a time t_0 when $\mathrm{ir}_\gamma(t_0)$ is realized by two different geodesics in $\overline{X}$. Then we can join the singularity at the endpoints of the geodesics in $B(\gamma(t_0), \mathrm{ir}_\gamma(t_0))$ and obtain a saddle connection as desired.

For every $t \in (0, l)$, we have a locally flat disk $B(\gamma(t), \mathrm{ir}_\gamma(t))$. We define the locally flat subset

$$B \coloneqq \bigcup_{t \in (0,l)} B(\gamma(t), \mathrm{ir}_\gamma(t)) \subseteq X.$$

By developing B into the plane we obtain an open, connected subset $\widetilde{B} \subseteq \mathbb{R}^2$. Every time we have σ on the boundary of $B \subseteq X$, we will consider a *representant* r of σ on the boundary of $\widetilde{B} \subseteq \mathbb{R}^2$. Note that the singularity σ is always the same in $\partial B \subseteq \overline{X}$ while the representants will not be identified in $\partial \widetilde{B} \subseteq \mathbb{R}^2$. Furthermore, let $\widetilde{\gamma}$ be the curve in $\widetilde{B} \subseteq \mathbb{R}^2$ which corresponds to the curve γ in B.

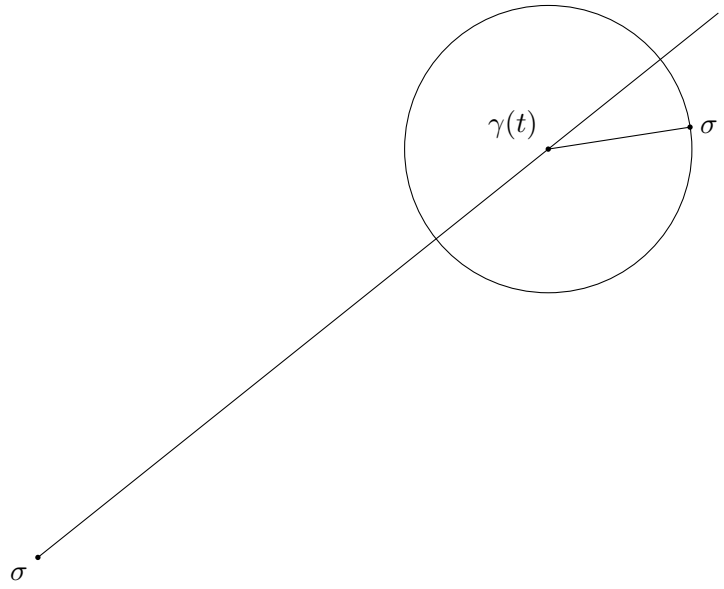

Figure 6: An example of a linear approach $[\gamma]$ and a time t so that $\mathrm{ir}_\gamma(t) < t$.

Define the set $R \coloneqq \{r \in \partial\widetilde{B} : r \text{ is a representant of } \sigma\} \subseteq \mathbb{R}^2$. This is a closed set in $\mathbb{R}^2$ as the corresponding point in $\overline{X}$ is always the same, namely σ.

For every representant $r \in R$ we can define the set

$$T_r \coloneqq \{t \in (0, l) : d(\widetilde{\gamma}(t), r) = \mathrm{ir}_\gamma(t)\}.$$

For every $r \in R$, the set T_r is closed in $(0, l)$ and connected:

- Let $(t_n)_{n\in\mathbb{N}}$ be a sequence in T_r converging to a time $t \in (0, l)$. We have
$$d(\widetilde{\gamma}(t_n), r) \leq d(\widetilde{\gamma}(t), r) + d(\widetilde{\gamma}(t), \widetilde{\gamma}(t_n)) = d(\widetilde{\gamma}(t), r) + |t - t_n|$$
and
$$d(\widetilde{\gamma}(t_n), r) \geq d(\widetilde{\gamma}(t), r) - d(\widetilde{\gamma}(t), \widetilde{\gamma}(t_n)) = d(\widetilde{\gamma}(t), r) - |t - t_n|$$
As ir_γ is continuous we can deduce
$$\mathrm{ir}_\gamma(t) = \lim_{n\to\infty} \mathrm{ir}_\gamma(t_n) = \lim_{n\to\infty} d(\widetilde{\gamma}(t_n), r) = d(\widetilde{\gamma}(t), r).$$
So t is in T_r and therefore T_r is closed in $(0, l)$.
- For the connectedness consider t_1, t_2 in T_r and $t_3 \in (0, l)$ such that $t_1 < t_3 < t_2$. Then for any $r' \in R$, we can compare the triangles in $\mathbb{R}^2$ with the vertices $\widetilde{\gamma}(t_1), \widetilde{\gamma}(t_2), r$ and $\widetilde{\gamma}(t_1), \widetilde{\gamma}(t_2), r'$, respectively (see Figure 7). As $d(\widetilde{\gamma}(t_1), r) \leq d(\widetilde{\gamma}(t_1), r')$ and $d(\widetilde{\gamma}(t_2), r) \leq d(\widetilde{\gamma}(t_2), r')$, we also have $d(\widetilde{\gamma}(t_3), r) \leq d(\widetilde{\gamma}(t_3), r')$ by elementary geometric arguments. So $t_3 \in T_r$.

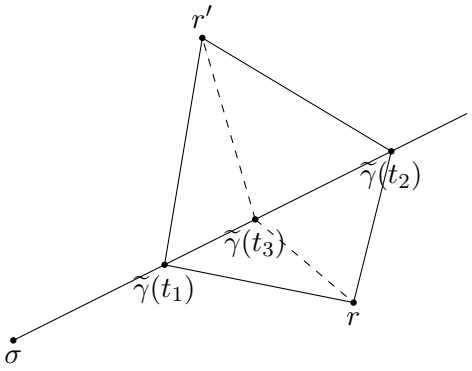

Figure 7: It the outer geodesics in one triangle are longer than the corresponding ones in the other triangle, then the third geodesic is longer, too.

We will now continue with a case-by-case analysis of $\partial\widetilde{B}$ and how it contains R:

Case 1: There is a connected subset of $\partial\widetilde{B}$ which is disjoint from R. Then there exists a closed connected subset b of $\partial\widetilde{B}$ whose interior is disjoint from R but whose endpoints (in a relative sense) are contained in R. We call these endpoints r_1 and r_2.

For technical reasons, we will now consider halfdisks instead of disks $B(\gamma(t), \mathrm{ir}_\gamma(t))$ and abuse notation. This means, we only consider the connected components of $B(\widetilde{\gamma}(t), \mathrm{ir}_\gamma(t))\setminus \mathrm{im}(\widetilde{\gamma})$ which are on the same side of γ as r_1 and r_2. Also, we only consider representants in R on the same side as r_1 and r_2, in T_r we consider times t where the geodesic from $\widetilde{\gamma}(t)$ to r is the shortest on the correct side, $\mathrm{ir}_\gamma(t)$ is the minimum of lengths of geodesics on the correct side, and so on.

In this sense we have that $T_{r_1}\cup T_{r_2}$ is connected: Choose $t_1 \in T_{r_1}$, $t_2 \in T_{r_2}$, and assume $t_1 < t_2$ (if not, change notation). Furthermore, choose $t_3 \in (0, l)$ with $t_1 < t < t_2$. Then we have for every $r \in R \setminus \{r_1, r_2\}$ on the same side of $\widetilde{\gamma}$ as r_1, r_2 that the inequality

$$d(\widetilde{\gamma}(t_3), r) \geq \min\{d(\widetilde{\gamma}(t_3), r_1),\, d(\widetilde{\gamma}(t_3), r_2)\}$$

holds, again by elementary geometric arguments. This means $t_3 \in T_{r_1}\cup T_{r_2}$ and $T_{r_1}\cup T_{r_2}$ is connected.

As T_{r_1} and T_{r_2} are closed and connected there exists a point $t_0 \in T_{r_1}\cap T_{r_2}$. In particular, the geodesics from $\widetilde{\gamma}(t_0)$ to r_1 and from $\widetilde{\gamma}(t_0)$ to r_2 in $\widetilde{B} \subseteq \mathbb{R}^2$ have the same length and in $B \subseteq X$ their interiors are both contained in $B(\gamma(t_0), \mathrm{ir}_\gamma(t_0))$. So we can join the endpoints of the two geodesics in $\partial B \subseteq \overline{X}$ and obtain a saddle connection from σ to σ of length less than $2 \cdot \mathrm{ir}_\gamma(t_0) \leq 2 \cdot t_0 \leq \epsilon$.

The obtained saddle connection is a chord of the halfdisk $B(\gamma(t_0), \mathrm{ir}_\gamma(t_0))$, so the corresponding linear approaches are contained in rotational components of length at least π. In fact, as there are no rotational components of length exactly π, the lengths have to be greater than π.

Case 2: Every connected subset of $\partial\widetilde{B}$ contains a representant r. This means, R is

dense in $\partial\widetilde{B}$. As R is a closed set, we can deduce $R = \partial\widetilde{B}$. Then we can choose two Cauchy sequences converging to different points in R. The distance of the corresponding elements of the sequence has to be bounded away from 0 in $\widetilde{B}$. Because of $R = \partial\widetilde{B}$ we have that the distance is also bounded away from 0 in X, so in $\overline{X}$ the two limit points are different singularities. This is a contradiction as we only have one singularity in $\partial B \subseteq \overline{X}$.

As case 2 can never happen and in case 1, we find a saddle connection as desired, the statement is proven. □

Theorem 1 (Wild singularity implies infinite genus)
Let (X, ω) be a translation surface with the following properties:

(i) The singularities of (X, ω) are discrete.

(ii) There exists a wild singularity σ such that none rotational component of σ has length π.

(iii) There exist two directions θ_1, θ_2 for which the geodesic flow $F_{\theta_1}, F_{\theta_2}$ is recurrent for almost every point in X.

Then X has infinite genus.

Proof. According to the criterion in Proposition 15 we want to show the statement by finding n saddle connections from σ to itself that intersect exactly in σ and so that the set has left-to-right-curves for every $n \in \mathbb{N}$. We do this by induction.

For $n = 1$ we can choose a saddle connection γ_1 so that both corresponding linear approaches are contained in a rotational component of length greater than π. Such a saddle connection exists by Proposition 21' and by Lemma 14 it is non-separating. Because of the careful choice of γ_1 we can define the immersion radius of γ_1 as in Definition 17 (iii).

So assume we have a non-separating set of $n-1$ saddle connections $\gamma_1, \dots, \gamma_{n-1}$ for which the immersion radii are defined. This means that for every $i \in \{1, \dots, n-1\}$ there exists a left-to-right curve δ_i, i.e. a curve in X that connects the left and the right side of γ_i without intersecting any of the γ_j, $j \in \{1, \dots, n-1\}$. Let ϵ be the minimum of all immersion radii of $\gamma_1, \dots, \gamma_{n-1}, \delta_1, \dots, \delta_{n-1}$. As proved in Proposition 21 there exists a saddle connection γ_n with length less than ϵ. Therefore, γ_n does not intersect $\gamma_1, \dots, \gamma_{n-1}, \delta_1, \dots, \delta_{n-1}$. Additionally, we can choose γ_n as in Proposition 21' so that both corresponding linear approaches are contained in a rotational component of length greater than π. Because of this, we can define the immersion radius of γ_n.

As γ_n is non-separating in X (see again Lemma 14) there exists a left-to-right curve δ_n' in $X \setminus \operatorname{im}(\gamma_n)$. If δ_n' has no intersection with $\gamma_1, \dots, \gamma_{n-1}$, we can define $\delta_n := \delta_n'$ and δ_n connects the left side of γ_n to the right side of γ_n in $X \setminus (\operatorname{im}(\gamma_1) \cup \dots \cup \operatorname{im}(\gamma_n))$. Furthermore, none of the curves $\delta_1, \dots, \delta_{n-1}$ intersects γ_n. Therefore we have that $\overline{X} \setminus (\operatorname{im}(\gamma_1) \cup \dots \cup \operatorname{im}(\gamma_n))$ is connected and the set of curves $\gamma_1, \dots, \gamma_n$ is non-separating.

If δ_n' intersects at least one of the curves $\gamma_1, \dots, \gamma_{n-1}$, we modify it in this way: For every intersection with a curve γ_i (without loss of generality, from the left of γ_i) we choose a point x_l on the left and a point x_r on the right of γ_i in $\operatorname{im}(\delta_n')$. Then we can replace the subcurve of δ_n' that intersects γ_i by a curve from x_l to the starting point of

δ_i concatenated with δ_i and concatenated with a curve from the endpoint of δ_i to x_r. By induction hypothesis and by the choice of ϵ, all δ_i ($i \in \{1, \dots, n-1\}$) do not intersect any of the curves $\gamma_1, \dots, \gamma_n$. Therefore the new curve δ_n is a left-to-right curve of γ_n with respect to $\gamma_1, \dots, \gamma_{n-1}$. Finally, we have that X has infinite genus. □

4 Possible generalizations of the theorem

One of the key points of the proof of Theorem 1 is the assumption of recurrence. By Poincaré recurrence one can deduce recurrence from the two weaker conditions that the flow is defined and the area is finite (see for example [HK02, Theorem 3.4.1]). We want to show in this section that none of the two conditions on its own works for the proof.

We start with an unpopular example that shows that recurrence is needed for the statement in Theorem 1:

Example 22 (Open disk). Let $X \subseteq \mathbb{R}^2$ be an open disk. We can endow X with the canonical translation structure ω that X carries as a subset of $\mathbb{R}^2$.

Then the metric completion $\overline{X}$ is the closed disk in $\mathbb{R}^2$. Therefore, the singularities of (X, ω) form an uncountable set that is in no way discrete.

For this translation surface (X, ω), there exists no direction and no point so that the point is recurrent under the geodesic flow in that direction. Not only are the methods of the proof of Theorem 1 not applicable but moreover it is obvious that X has genus 0.

The open disk obviously doesn't fulfill the recurrence condition as it has a whole continuum of singularities that is "trapping" all rays. Now we can ask if the conditions of discrete singularities and finite area are sufficient to deduce recurrence. For this, we consider the following example:

Example 23 (Icicled surface). Consider a half-open rectangle of height 2 and width 1. The left side is glued to the right side, the bottom and the top are not included.

For every $n \in \mathbb{N}$, we consider a vertical segment starting at the bottom, resp. top at $\frac{i}{2^n}$ of length $\frac{1}{2^n}$ for every odd $i \in \{1, \dots, 2^n - 1\}$ (see Figure 8). We call this vertical segments *icicles*.

We glue the segments as sketched in Figure 9. Note that no icicle on the top is glued to an icicle on the bottom. Formally, we can describe the gluings for the left part of the top by:

- For each side of the icicle at $\frac{1}{2}$, we cut the segment again: First we cut it in half, then we cut the upper half into halfs again, cut the upper quarter into halfs again, ... So for every $n > 1$ we have a segment of length $\frac{1}{2^n}$.
- The left side of the lower half of the icicle at $\frac{1}{2}$ is glued to the right side of the icicle at $\frac{1}{4}$.
- For every $n > 2$ and every odd $i \in \{3, \dots, 2^{n-1} - 1\}$, the left side of the icicle at $\frac{i}{2^n}$ is glued to the right side of the icicle at $\frac{i-2}{2^n}$.

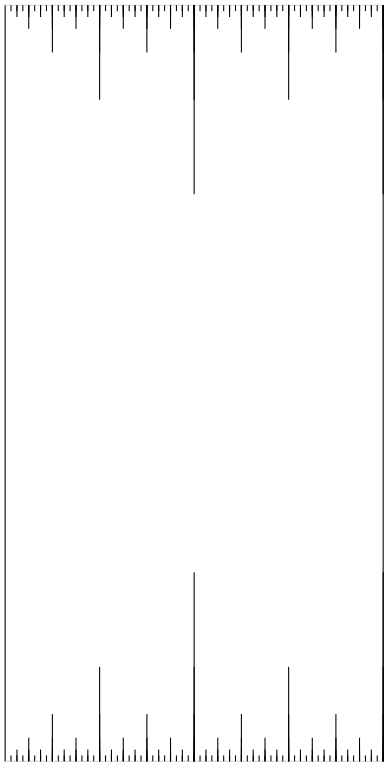

Figure 8: Icicled surface.

- For every $n > 2$, the left side of the icicle at $\frac{1}{2^{n-1}}$ is cut into two segments of the same length. The lower part is glued to the right side of the icicle at $\frac{2^{n-1}-1}{2^n}$. The upper part is glued to the right side of the segment at $\frac{1}{2}$ which has the correct length.
- We do the similar gluing for the right part of the top and also for the bottom.

Then (X, ω) has the following properties:

(i) The icicled surface X endowed with the canonical translation structure ω is a translation surface as every point in X has a flat neighborhood.

(ii) The translation surface (X, ω) has two singularities: All the tips of the icicles on the top are identified by the definition of the gluings. We call the corresponding singularity σ_{top}. Now, we consider a non-dyadic point p in the top boundary of the rectangle, i.e. a point where no icicle starts. There exists a sequence of icicles such that the tips of the icicles converge to p, seen as points in $\mathbb{R}^2$ without gluings. Therefore, the distance from σ_{top} to p is 0 and $p = \sigma_{top}$ in $\overline{X}$. The same argument works for the points where we cut some icicles into more segments and for the points on the boundary where an icicle starts.

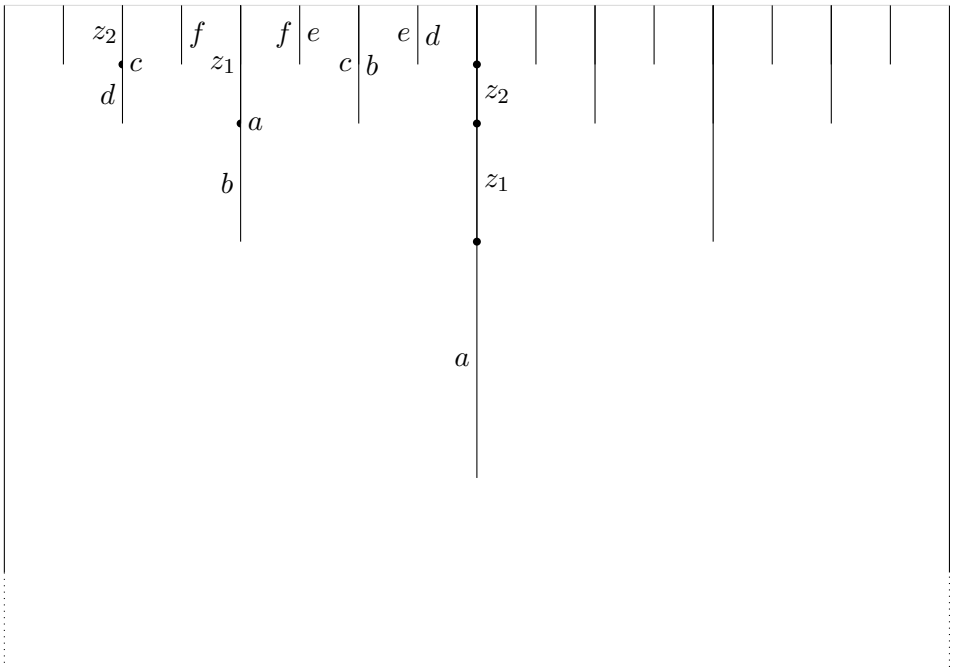

Figure 9: Gluings for the icicled surface: segments with the same letters are glued.

Also, the same reasoning holds for the icicles on the bottom and the bottom boundary. So we have only two singularities σ_{top} and σ_{bottom}.

For the record, both of the singularities have uncountably many rotational components of finite length but none of the rotational components has length π.

(iii) The surface X has infinite genus: Here we cannot use the arguments of Theorem 1 (see next item). However, we can check by a sharp look that every icicle is defining at least one saddle connection which is non-separating. Moreover, every set of saddle connections defined by icicles is non-separating. As the number of icicles is not bounded, we can show according to Proposition 15 that X has infinite genus.

(iv) Apparently, the geodesic flow in the vertical direction is not defined for any point in X for all time.

Let θ be a direction in $(0, \pi)$. Suppose, the geodesic flow F_θ is defined for all time for a set of positive measure. Consider a closed horizontal geodesic g in the middle of the surface. Then F_θ is defined for a subset of g of positive measure for all time. From Poincaré recurrence we can deduce that there is a point $x \in g$ and a time t_x so that $F_\theta(x, t_x) \in g$. This means that $F_\theta(x, t_x - \epsilon)$ for an $\epsilon > 0$ is contained in the lower part of the surface. But this is impossible as there is no possibility to reach the lower part of the surface from the upper part of the surface.

So for all directions except of the horizontal one the geodesic flow is not defined for

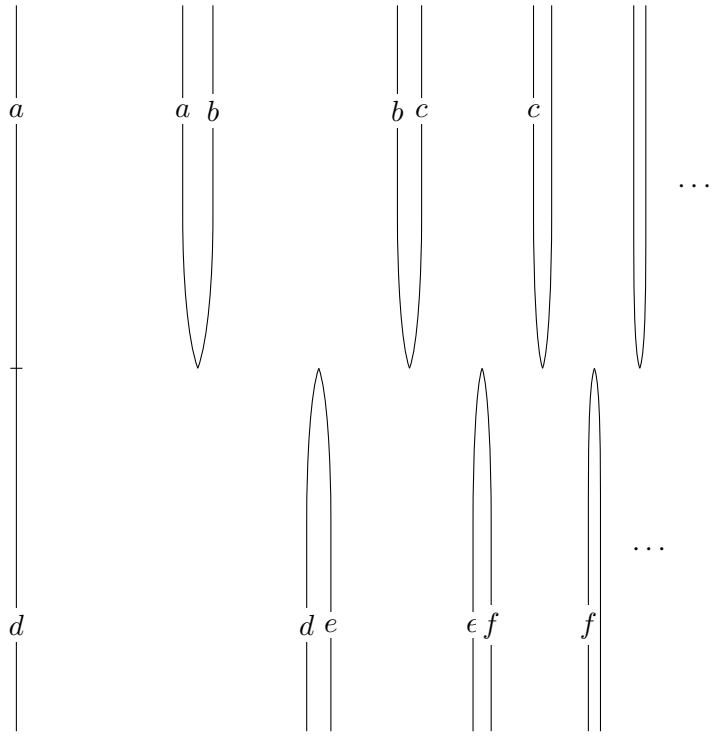

Figure 10: Translation surface with a wild singularity and genus 0.

a set of positive measure for all time.

(v) The horizontal flow is periodic but the length of the period is getting greater the nearer to the top or to the bottom the starting point is.

The icicled surface example shows the following negative result:

Remark 24 (Discrete singularities do not necessarily imply good dynamical properties). There exist translation surfaces with discrete singularities so that for at most one direction θ the geodesic flow F_θ is defined for almost every point for all time.

The converse statement on the flow is needed as an assumption in Lemma 14. In fact, there are translation surfaces with discrete singularities and separating saddle connections. For example, the horizontal saddle connection of the icicled surface in Example 23 that connects the tip of the longest icicle to itself is separating.

The next example was shown to the author by Pat Hooper. It will illustrate that it is not sufficient for recurrence, that for two directions the geodesic flow is defined for almost every point for all time:

Example 25 (Nested cylinders). Consider a Euclidean half plane with a distinguished midline. We cut vertical slits of infinite length in the half plane from the midline upward,

starting from

$$\frac{1}{2} + \sum_{i=2}^{n} \frac{1}{2i-1} + \frac{1}{2i} \quad \text{for } n \geq 1.$$

Additionally, we cut vertical slits of infinite length from the midline downward, starting from

$$\sum_{i=1}^{n} \frac{1}{2i} + \frac{1}{2i+1} \quad \text{for } n \geq 1.$$

Now we glue the right side of a slit to the left side of the slit which is next to the right and the left side of the slit to the right side of the slit which is next on the left (see Figure 10). By this construction, we obtain half-infinite cylinders with smaller and smaller circumferences that are glued in a nested way.

The resulting translation surface has infinite area, genus 0 and exactly one singularity. The singularity is wild as the distance between the startings points of the slits is going to zero, i.e. there exists no cyclic translation covering from a neighborhood of σ to a neighborhood of $0 \in \mathbb{R}^2$. This singularity has exactly one rotational component which is of infinite length.

For every direction θ, the geodesic flow F_θ is defined for almost every point for all time. However, recurrence only occurs in the horizontal direction.

In particular, all saddle connections are horizontal and all of them are separating.

The last example shows that in Theorem 1 assumption (iii) can not be dropped without replacement. For example, a replacement by the condition of finite area is not excluded by the shown examples although an adaption of the proof would be needed.

The discreteness of the singularities in assumption (i) and the existence of a wild singularity in (ii) is crucial for the proof given here. However, the condition on the length of rotational components in assumption (ii) arises merely by technical details in the method, more precisely in Proposition 21'.

References

[AH12] Artur Avila and Pascal Hubert, *Recurrence for the wind-tree model*, preprint (2012).

[Bow12] Joshua P. Bowman, *Finiteness conditions on translation surfaces*, Providence, RI: American Mathematical Society (AMS), 2012.

[Bow13] ———, *The complete family of Arnoux-Yoccoz surfaces*, Geometriae Dedicata **164** (2013), 113–130.

[BV13] Joshua P. Bowman and Ferrán Valdez, *Wild singularities of flat surfaces*, Israel Journal of Mathematics **197** (2013), 69–97.

[Cha04] Reza Chamanara, *Affine automorphism groups of surfaces of infinite type*, In the Tradition of Ahlfors and Bers, III (William Abikoff and Andrew Haas, eds.), Contemporary mathematics, vol. 355, 2004, pp. 123–145.

[DHL11] Vincent Delecroix, Pascal Hubert, and Samuel Lelièvre, *Diffusion for the periodic wind-tree model*, preprint (2011).

[FK36] Ralph H. Fox and Richard B. Kershner, *Concerning the transitive properties of geodesics on a rational polyhedron*, Duke Mathematical Journal **2** (1936), 147–150.

[FU14] Krzysztof Frączek and Corinna Ulcigrai, *Non-ergodic $\mathbb{Z}$-periodic billiards and infinite translation surfaces*, Inventiones mathematicae **197** (2014), no. 2, 241–298.

[Ful97] William Fulton, *Algebraic topology : a first course*, corr. 2. print. ed., Graduate texts in mathematics, vol. 153, Springer, New York, 1997.

[Ghy95] Etienne Ghys, *Topologie des feuilles génériques*, Annals of Mathematics. Second Series **141** (1995), no. 2, 387–422.

[HHW13] W. Patrick Hooper, Pascal Hubert, and Barak Weiss, *Dynamics on the infinite staircase*, Discrete and Continuous Dynamical Systems **33** (2013), no. 9, 4341–4347.

[HK02] Boris Hasselblatt and Anatole Katok, *Principal structures*, Handbook of dynamical systems. Volume 1A, Amsterdam: North-Holland, 2002, pp. 1–203 (English).

[HLT11] Pascal Hubert, Samuel Lelièvre, and Serge Troubetzkoy, *The Ehrenfest wind-tree model: periodic directions, recurrence, diffusion*, Journal für die Reine und Angewandte Mathematik **656** (2011), 223–244.

[Hoo10] W. Patrick Hooper, *The invariant measures of some infinite interval exchange maps*, preprint (2010).

[Hoo14] ———, *An infinite surface with the lattice property I: Veech groups and coding geodesics*, Transactions of the American Mathematical Society **366** (2014), no. 5, 2625–2649.

[HS10] Pascal Hubert and Gabriela Schmithüsen, *Infinite translation surfaces with infinitely generated Veech groups*, Journal of Modern Dynamics **4** (2010), no. 4, 715–732.

[HW12] W. Patrick Hooper and Barak Weiss, *Generalized staircases: recurrence and symmetry*, Annales de l'Institut Fourier **62** (2012), no. 4, 1581–1600.

[LT14] Kathryn Lindsey and Rodrigo Treviño, *Flat surface models of ergodic systems*, preprint (2014).

[PSV11] Piotr Przytycki, Gabriela Schmithüsen, and Ferrán Valdez, *Veech groups of Loch Ness monsters*, Annales de l'Institut Fourier **61** (2011), no. 2, 673–687.

[Ric63] Ian Richards, *On the classification of noncompact surfaces*, Transactions of the American Mathematical Society **106** (1963), no. 2, 259–269.

[Tre14] Rodrigo Treviño, *On the ergodicity of flat surfaces of finite area*, Geometric and Functional Analysis **24** (2014), no. 1, 360–386.

[Vee89] William A. Veech, *Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards*, Inventiones Mathematicae **97** (1989), no. 3, 553–583.